

Automatic Mapping of Linear Woody Vegetation Features in Agricultural Landscapes Using Very High-Resolution Imagery

Selim Aksoy, *Member, IEEE*, H. Gökhan Akçay, Tom Wassenaar

Abstract—Automatic mapping and monitoring of agricultural landscapes using remotely sensed imagery has been an important research problem. This paper describes our work on developing automatic methods for the detection of target landscape features in very high spatial resolution images. The target objects of interest consist of linear strips of woody vegetation that include hedgerows and riparian vegetation that are important elements of the landscape ecology and biodiversity. The proposed framework exploits the spectral, textural, and shape properties of objects using hierarchical feature extraction and decision making steps. First, a multi-feature and multi-scale strategy is used to be able to cover different characteristics of these objects in a wide range of landscapes. Discriminant functions trained on combinations of spectral and textural features are used to select the pixels that may belong to candidate objects. Then, a shape analysis step employs morphological top-hat transforms to locate the woody vegetation areas that fall within the width limits of an acceptable object, and a skeletonization and iterative least-squares fitting procedure quantifies the linearity of the objects using the uniformity of the estimated radii along the skeleton points. Extensive experiments using QuickBird imagery from three European Union member states show that the proposed algorithms provide good localization of the target objects in a wide range of landscapes with very different characteristics.

Index Terms—Multi-scale texture analysis, shape analysis, linear object detection, object-based performance evaluation.

I. INTRODUCTION

Among all economic sectors, agriculture is by far the largest land user and is increasingly being seen as a potential major steward of our global ecosystem, which adds an important dimension to its traditional production goals. The European Union (EU), upon gradually transforming its Common Agricultural Policy (CAP) toward a more liberalized system, has built into the CAP a number of environmental safeguards aiming at avoiding negative impacts of deregulation, as well as reducing the current environmental impact of the sector. The detailed definition of these safeguards, so-called Good Agricultural and Environmental Conditions (GAECs), part of cross-compliance standards to be respected by farmers claiming subsidies, are set by the EU Member States. Some

member states have defined rules obliging farmers to play an active role in landscape and habitat maintenance, like e.g., the maintenance of hedgerows and riparian vegetation belonging to the farm property that constitute an important element of the landscape's ecological infrastructure. Adoption of such standards and regulations at regional, national, and international levels require active monitoring of their enforcement.

Remote sensing has long been acknowledged as an important tool for planning and monitoring of land cover/use. Development of automatic and robust methods has become an important research problem when the analysis goes beyond local sites to cover a wide range of landscapes in national and even international levels. Such methods are also paramount when monitoring change over time. In order for remote sensing to allow for the efficient enforcement monitoring of the aforementioned regulations, the goal of this study is to develop automatic methods for detailed mapping of target landscape features in very high spatial resolution images.

The target objects of interest in this paper are *linear strips of woody vegetation* separating agricultural fields with examples shown in Figure 1. These objects include hedges/hedgerows defined as a row of bushes or trees planted closely to form a boundary between pieces of land or at the sides of a road, and riparian vegetation defined as a narrow zone with woody plant communities along river or stream margins, bordered on the other side by agricultural land [2]. They are important biological and ecological components of the environment where they serve many functions including providing field boundaries, animal habitats, windbreaks, erosion control, and contributing to landscape ecology and biodiversity [3], [4]. In the EU, several member states have as of 2007 explicitly included the maintenance of such landscape features among their GAEC standards. These include Cyprus, Czech Republic, France, Germany, Luxembourg and the United Kingdom (see [5] for the UK). However, none of the corresponding control authorities dispose as yet of an efficient monitoring tool.

Classification of land cover has traditionally been performed using pixel level processing with mainly statistical tools in a multi-class setting. These multi-class classifiers need example patterns for each class to estimate decision boundaries in the feature space during training. However, in real world classification problems, sampling a sufficient number of training data from each of the classes is not always possible. Since these classifiers require complete descriptions of all classes, they may not generalize well with a sufficiently high accuracy for a large number of classes, especially when some of them

Manuscript received February 19, 2009; revised April 23, 2009. This work was supported in part by the European Commission Joint Research Centre contract 253352 and TUBITAK CAREER grant 104E074. An early version of this paper was presented in IGARSS 2008 as [1].

S. Aksoy and H. G. Akçay are with the Department of Computer Engineering, Bilkent University, Ankara, 06800, Turkey. Email: {saksoy,akçay}@cs.bilkent.edu.tr.

T. Wassenaar is with the European Commission, Joint Research Centre, Institute for the Protection and Security of the Citizen, Via E. Fermi 2749, TP 266, 21027 Ispra (VA), Italy. Email: tom.wassenaar@jrc.it.




Fig. 1. Example QuickBird images (pan-sharpened visible bands) containing linear strips of woody vegetation marked with a yellow boundary by an expert. *Linearity* of their shape and *woodiness* of their texture distinguish these structures from the rest of the image. Linearity is defined here as piecewise elongation along the major axis while having an approximately constant width, not necessarily in the strict sense of a perfectly straight line. Woody vegetation consists of trees and bushes. Note that other linear but non-woody structures and woody vegetation areas that are not linear are not marked. (Raster images in this paper are 1000×1000 pixels in size corresponding to 600×600 m.)

have large variations in appearance. Therefore, delineation of individual trees or tree groups is not necessarily very accurate when the goal is to classify the whole land cover. Generic object-based classification is also not suitable here because a holistic analysis requires an image-wide prior segmentation, but accurate segmentation of very high spatial resolution images is still a very hard problem [6]. Furthermore, detailed a priori information about the shapes of the objects of interest cannot be easily incorporated into the classification process in a multi-class setting.

Alternatively, one can set up the problem as the detection of single, pre-defined objects where the methods concentrate on specific properties of these objects. An example for such detection is the identification of vegetation by thresholding specific spectral features such as the normalized difference vegetation index, soil-adjusted vegetation index or atmospheric resistant vegetation index [7], [8]. However, these indices consider only the spectral properties of individual pixels, and do not take into account spatial and contextual information.

Another method that is widely used for detecting pre-defined objects is template matching. Templates are often defined manually or are learned from examples provided by the user, and detection is performed by moving the template over the image and evaluating the match at each location using a similarity measure such as correlation [9]. However, these templates are often fixed in terms of size, shape, and intensity, and cause the detection algorithm to have problems regarding invariance to scale, rotation, and illumination.

Quackenbush [9] published a review of techniques for linear feature detection in images. Popular techniques include mathematical morphology, Hough transform, multi-resolution edge detection, template matching, dynamic programming for

edge linking, and rule-based classification. Such techniques have been applied to the extraction of roads [10], buildings [11], and water channels [12]. However, they are not directly applicable to the detection of linear strips of woody vegetation because they assume the existence of collinear and parallel line segments that constitute pairs of edges forming object boundaries, whereas the textured agricultural regions often produce a lot of small line segments both within and along the boundaries, and edges that can be detected along the boundaries of vegetation regions also show a lot of irregularities. Furthermore, rural linear features such as hedgerows and riparian vegetation often exhibit directional variation according to whether they follow natural boundaries such as streams and rivers or human-made linear objects such as roads, or they have been planted as separators between agricultural fields [13].

Several studies for the analysis of hedgerows have been published in the literature. Most of these studies concentrate on the functional categorization of hedgerows and their development in time where the mapping is already known or is done by manual photo-interpretation [3], [14]. As an example for automatic detection, Lennon *et al.* [15] used a fuzzy combination of a vegetation index, a linearity feature based on image gradient, and co-occurrence texture features for data obtained from an airborne hyperspectral sensor. However, no experiments were presented so the accuracy of the above method is unknown. Thornton *et al.* [4] used a resolution enhancement algorithm for sub-pixel mapping of linear thin structures such as hedgerows with 1-3 m width in 10 m Spot images. The technique was illustrated on a small image that contained a single hedgerow. More recently, Vannier and Hubert-Moy [16] compared the usefulness of an orthophotoplan, a Spot 5 image, an Aster image, and a Landsat image for detecting hedgerows. The eCognition software was used for segmenting each image by adjusting the parameters such as object size, color homogeneity, shape, smoothness, and compactness individually so that segmentations were obtained at three scales corresponding to tree, hedge, and field levels. Different spectral features and fuzzy memberships functions were used to classify the resulting objects. Tansey *et al.* used a similar approach in [17]. However, the need for individually adjusting many different parameter values for multi-scale segmentation and classification for each image may limit the general applicability of this method.

The detailed content of very high-resolution imagery and the large spatial coverage of such data sets require the development of new techniques for detection of individual pre-defined objects. Our main contributions toward the detection of linear strips of woody vegetation (called *hedges* in the rest of the paper) as the target objects of interest are twofold. The first is a framework that exploits the spectral, textural, and object shape information using hierarchical feature extraction and decision making steps. First, pixel-based spectral and multi-scale textural features are extracted from the input panchromatic and multispectral data. Then, discriminant functions trained on combinations of these features are used to select the pixels that may belong to targets of interest, and connected components analysis on these pixels is used to obtain the candidate objects (woody vegetation). A multi-feature and multi-scale strategy

was used to be able to cover different textural characteristics generated by individual trees and their different groupings and appearances in a wide range of sites. Shackelford and Davis [18] also showed the effectiveness of a multi-feature hierarchical strategy in the classification of high-resolution images where the spectral bands were initially used to split the data into grass-tree, road-building, water-shadow, and bare soil classes, and then the entropy texture feature was used for further separation of grass-tree whereas the length-width feature was used for road-building and water-shadow classes.

Our second main contribution is a novel algorithm that can identify the linear structures within the candidate areas, and can even separate the target objects of interest from other tree groups. This algorithm involves morphological top-hat transforms to locate the woody vegetation areas that fall within the width limits of an acceptable hedge, and a skeletonization and iterative least-squares fitting procedure that quantifies the linearity of the objects using the uniformity of the estimated radii along the skeleton points. The width and length criteria of structures can easily be adjusted according to landscape characteristics thereby also allowing for the separate detection of different object classes or class instances. The parts of the candidate objects that pass these tests are labeled as detected targets (hedges). This shape model is generic enough such that it can be adapted to the detection of other linear object classes with natural boundaries instead of strictly straight boundaries (e.g., rivers, roads, and paths that are bordering natural land cover classes). Extensive experiments using QuickBird imagery from three European sites with different characteristics show that the proposed algorithms provide good localization of the target objects in different landscapes.

The rest of the paper is organized as follows. The study sites and the corresponding data are described in Section II. Pixel-based spectral and textural feature extraction is presented in Section III. Identification of woody vegetation areas as candidate objects is discussed in Section IV. Object-based shape feature extraction for the quantification of the linearity of these objects and the final decision for target detection based on the extracted shape features are described in Section V. Experiments using ground truth data and the details of the object-based quantitative performance measures are presented in Section VI. Finally, conclusions are given in Section VII.

II. STUDY SITES AND DATA

The annual Control With Remote Sensing (CWRS) program of the EU CAP is among the world's largest civil programs where remote sensing data are employed on a regular basis. It is noteworthy that, to date, very high spatial resolution (VHR) imagery accounts for more than half of the CWRS budget for remote sensing data acquisition. IKONOS and QuickBird are by far the most used VHR sensors in the CWRS program. Panchromatic and pan-sharpened QuickBird-2 sensor data with 60 cm spatial resolution were employed in this study, considered to better represent characteristics of near future VHR satellite imagery. The data used were from three EU member states with a hedge conservation GAEC standard.

Performance evaluation was done using 33 subscenes, each with size 1000×1000 pixels, cut from three QuickBird

images of Baden-Württemberg, Germany; Decin, Czech Republic; and Paphos, Cyprus (11 subscenes were used from each image). These sites were chosen to collect a diverse sample of hedges with different characteristics. The Baden-Württemberg site is a rolling agricultural landscape typical of large parts of the temperate EU, with large clumps of variably sized agricultural parcels intersticed with medium and large forest patches. Hedges are nearly exclusively parcel separations. Pasture dominated Decin site hedges are much larger on average and riparian vegetation is more frequent. Paphos site represents a rather extreme situation of thin hedges in a very fragmented environment containing many other small linear features. Examples are shown in Figure 1. The sites are referred to as *Baden*, *Decin*, and *Paphos*, respectively, in the rest of the paper.

III. FEATURE EXTRACTION

Spectral features can be used to distinguish green vegetation from the rest of the image. Texture features are useful for identifying areas that have similar spectral responses but different spatial structures. An important consideration in this work was that the desired features not only could describe image windows but were also able to localize the structures of interest within these windows. Even though features that are based on first-order statistics (e.g., mean, variance, skewness) and second-order statistics (e.g., co-occurrence features) of spectral values have been shown to be effective in characterizing texture in rectangular image windows, they were not suitable for the problem in this work because they could not necessarily localize specific structures within these windows.

We observed that two different types of textural characteristics were important: the arrangements of individual trees and the appearance of linear structures with respect to their surroundings. Therefore, multi-scale texture features were considered.

A. Normalized difference vegetation index

The normalized difference vegetation index (NDVI) [8] is a simple but powerful measure for identifying photosynthetically active vegetation. NDVI, computed from the pan-sharpened multispectral data, was used to separate green vegetation from the rest of the land cover. We observed that, although it might not distinguish hedges from other types of vegetation, it was useful for eliminating linear human-made structures that could cause false alarms in the decision process.

B. Gabor features

Gabor features have been very popular in the texture analysis literature due to their good localization abilities in both spatial and frequency domains with flexibility for multi-scale and multi-orientation tunability. Gabor features were extracted by applying a bank of scale and orientation selective filters [19] to the panchromatic band. In particular, 6 scales and 6 orientations were used with a resulting set of 36 bands. The scales used were designed to include both the fine texture of individual trees within a hedge and the coarse texture of hedges




Fig. 2. Example features for some of the images in Figure 1. Each row shows the features for one image. Feature values are scaled for better visualization.

among agricultural fields. To obtain rotation invariance, the responses for all filters with different orientations at a given scale were combined using the pixelwise “max” operator.

C. Granulometry features

The concept of granulometry is based on the notion of sieving a sample with increasing sieve sizes so that increasing number of grains fall through the sieve. The measured mass for each sieve size creates a size distribution, also called the pattern spectrum because its peaks indicate the prevailing sizes of the structures [20].

The concept of granulometry can be transposed to image data by morphological opening and closing of the image with a family of structuring elements with increasing sizes [20]. A granulometry by opening produces information concerning image structures brighter than their neighborhood. Similarly, granulometry by closing gives information about the arrangement of structures that appear darker than their neighborhood. Local granulometries can be computed by summing the pixel values within sliding windows after every opening and closing operation to represent the size distribution. We computed morphological granulometry features from the panchromatic band using opening and closing with disk structuring elements with radii from 1 to 9 pixels in steps of 2. This resulted in a set of 10 features. The scales (structuring element sizes) were selected by visual examination of the feature results.

Overall, NDVI was extracted from the pan-sharpened multispectral data as the spectral feature, and Gabor and granulometry features were extracted from the panchromatic data for modeling texture. Examples are shown in Figure 2.

IV. IDENTIFICATION OF CANDIDATE OBJECTS

After the features are extracted, the next step is to find the image areas that give high responses to these features so that they can be considered as candidate objects. We used a two-step decision process. First, a threshold on NDVI was used to separate green vegetation from the rest of the land cover. The threshold was selected so that there was no omission of any hedge structure. However, we observed that such thresholding could not distinguish hedges from other types of vegetation and kept many fields, large groups of trees and other vegetated areas in the output. On the other hand, the thresholding eliminated some linear human-made structures that gave high responses to the texture features.

Given the obtained vegetation mask, the next step is to identify candidate objects according to their texture characteristics. Pixel-based texture modeling was not sufficient for detecting the linearity of a structure but was capable of modeling its woodiness. Therefore, we concentrated on the separation of woody vegetation from the rest of the areas that appeared in the vegetation mask.

A. Woody vs. non-woody vegetation classification

Manual labeling of image areas as woody vs. non-woody vegetation was used to generate the ground truth for training and evaluation. A randomly selected subset consisting of 750,000 pixels was used for training and another independent subset consisting of 375,000 pixels was used for validation.

Different combinations of features were studied, and several classifiers such as Gaussian maximum likelihood classifier,

TABLE I

THE ACCURACY OF PIXEL LEVEL CLASSIFICATION FOR WOODY VS. NON-WOODY VEGETATION USING DIFFERENT FEATURE COMBINATIONS. TP: PERCENTAGE OF TRUE POSITIVES, TN: PERCENTAGE OF TRUE NEGATIVES, AVG: AVERAGE ACCURACY.

MS	Gabor	Gran.	Selection	TP	TN	Avg.
X				84.88	81.99	82.95
	X			73.03	84.71	80.82
		X		75.54	87.52	83.52
X	X			91.22	91.56	91.45
X		X		93.70	91.43	92.18
	X	X		78.97	87.56	84.70
X	X	X		93.40	91.71	92.27
X	X	X	X	93.92	96.65	94.83

mixture of Gaussian classifier, naive Bayes classifier, and linear Fisher classifier were compared. These classifiers were selected because of their simplicity as the goal was to use models that were as simple as possible so that they could be trained with as few examples as possible. More complex classifiers such as neural networks or support vector machines were not used because of the ambiguities in parameter selection and high computational complexity for large data sets.

We observed that there was no significant difference in the accuracies of different classifiers, and the Gaussian classifier performed as good as any other classifier, so the quadratic Gaussian maximum likelihood classifier was used in the rest of the paper.

B. Performance evaluation

We performed an exhaustive study to evaluate the effectiveness of different feature combinations. The features considered included 4 multispectral values, 6 Gabor features, and 10 granulometry features, resulting in a total of 20 features for each pixel. Table I shows the classification rates for different combinations. Among the features, combining multispectral bands with texture features performed better than using each type of features individually. However, comparing the combinations of a particular set of texture features with multispectral bands did not show any significant difference among different combinations. In general, we observed that some features performed better than others but using all features together did not provide any significant improvement. This suggested that there were correlation and redundancy among the features and using all of them did not provide any additional information. Therefore, we decided to perform feature selection to obtain a good subset of features.

We used the sequential backward selection algorithm [21] that is an iterative algorithm that starts with all features and shrinks down the feature set by, at each iteration, removing the single worst feature from the set of features obtained in the previous iteration. The results showed that feature selection actually improved the classification accuracy. The best result (94.83%) was obtained when the original set of 20 features were reduced to 9 features. These features consisted of all four multispectral values (blue, green, red, near infrared), two Gabor scales (4 and 6), and three granulometry scales (closing 1, opening 1, opening 5). The selection process favored both texture features and multispectral values. Even though

TABLE II

THE AVERAGE ACCURACY OF WOODY VS. NON-WOODY CLASSIFICATION FOR INDIVIDUAL SITES. THE ROWS CORRESPOND TO TRAINING DATA SOURCES. THE COLUMNS CORRESPOND TO VALIDATION DATA SOURCES.

	<i>Baden</i>	<i>Decin</i>	<i>Paphos</i>
<i>All</i>	95.87	97.30	84.91
<i>Baden</i>	96.41		
<i>Decin</i>		97.96	
<i>Paphos</i>			89.47

the multispectral values were not sufficient for identifying woody vegetation on their own, they helped with localization, especially at texture boundaries when larger texture scales (larger filter windows) were used, and helped producing more accurate boundaries along the objects of interest. We also evaluated feature reduction using principal components analysis, but this method did not give good results, as expected, because the resulting features were not optimal for discrimination.

We also compared the accuracy of woody vs. non-woody vegetation classification within individual sites. The results in Table I used the classifier trained with samples collected from all sites and validated with samples also collected from all sites. Therefore, these results were a good indicator of the cross-landscape performance of the classifier on sites with very different characteristics. We considered 6 experiments for 4 training scenarios for performance on individual sites:

- training using samples from all sites and validating using samples from individual sites,
- training using samples only from individual sites and validating using samples only from individual sites (3 scenarios).

The results in Table II showed that site specific classifiers could improve the accuracy even further. The most significant improvement was observed for the *Paphos* data set. This was due to the fact that the woody vegetation in the *Paphos* sites showed much different scale and texture characteristics than the vegetation in the *Baden* and *Decin* sites, and the classifier that could capture these characteristics from the site specific training data could result in a higher accuracy. The object detection experiments in Section VI use both the woody vegetation classifier trained using samples from all sites and the classifiers trained using site specific samples.

After the discriminant function identified the pixels that could belong to targets of interest, connected components labeling of these pixels was used to obtain the candidate objects. Morphological opening and closing operations were used to eliminate small noise components and to fill small holes. Example results are shown in Figure 3.

V. DETECTION OF TARGET OBJECTS

After the candidate objects are found, object shape information was used so that the objects can be labeled as target or are rejected. An important observation was that the results of the connected components labeling in the previous step were not directly suitable for computing object level features. The reasons were twofold: hedges were often connected to other larger groups of trees, and they often followed natural boundaries where they did not necessarily exhibit a perfectly




Fig. 3. Example results for woody vs. non-woody vegetation classification after morphological cleaning of the connected components obtained from the pixel level classification results. The image areas identified as woody vegetation are marked as green on the panchromatic image. Note that woody vegetation can have very different appearances in different sites.

straight structure. Therefore, commonly used shape features such as eccentricity, major/minor axes, orientation, and moments were not good indicators of the linearity of the shapes of these objects. Hence, an important step was the separation of hedges from other tree groups and piecewise linearization of the object regions before extracting the shape features.

The object-based feature extraction process used skeletonization as a structural representation of the object shapes, and morphological filtering and iterative least-squares fitting based segment selection to extract the parts of this representation that might correspond to a hedge. First, the skeleton of the binary classification map of candidate objects was computed. The output of this step was the set of points on the skeleton, and, for each point, the radius of the maximal disk that was centered at that point and was contained within the binary shape. However, skeletonization is known to be sensitive to small intrusions and extrusions so a pruning is often necessary before further processing. Therefore, we pruned the initial skeleton by recursively removing all branches that were shorter than a length threshold until no such branch remained. Example skeletons are shown in Figures 4(a) and 4(d).

Further pruning of the candidate objects was done using morphological filtering to find the objects that fall within the shape limits of an acceptable hedge. Given two thresholds that specified the maximum and minimum acceptable width of a hedge, first, a disk structuring element $S_{\max\text{-width}}$ with diameter slightly larger than the first threshold and another disk structuring element $S_{\min\text{-width}}$ with diameter slightly smaller than the second threshold were constructed. Then, potential hedges were extracted from the set of candidate objects using consecutive application of top-hat transforms and conditional dilations. The morphological top-hat transform was computed as the difference between the candidate object image I and its opening with a particular structuring element S as

$$\text{TH}(I, S) = I - (I \circ S). \quad (1)$$




Fig. 4. Example results for object-based feature extraction. The first column shows initial skeletons overlaid on the woody classification maps. The second column shows the parts that remained after morphological top-hat filtering. The third column shows the objects corresponding to the final set of segments selected as linear using the least-squares fitting procedure.

Hence, it removed the image structures that could contain the structuring element used. In particular, the first top-hat transform was computed using the larger structuring element and identified the structures that satisfied the maximum width requirement. The second top-hat transform was computed using the smaller structuring element and identified the structures that were narrower than an acceptable hedge. The structures that were in the result of the first top-hat but not in the result of the second top-hat were extracted as

$$I_{\text{new}} = \text{TH}(I, S_{\max\text{-width}}) - \text{TH}(I, S_{\min\text{-width}}) \quad (2)$$

as potential hedges, and only the parts of the skeleton that corresponded to these objects were kept for further analysis. Example results for morphological filtering are shown in Figures 4(b) and 4(e).

The morphological filtering step eliminated the structures that were too wide or too narrow. This also decreased the computation time by excluding the structures that were not within the shape limits of an acceptable hedge from further processing. However, it did not guarantee that the remaining structures were linear. The next step used iterative least-squares fitting based segment selection. First, lists of skeleton points that were connected to each other and were separated by junctions were found. The output of this step was a set of linked point lists representing the objects or segments of objects on the skeleton. We assumed that the linearity of a segment could be modeled by the uniformity of the radii along the skeleton points that corresponded to the uniformity of the width perpendicular to the medial axis. This assumption was implemented using a line fitting procedure that was applied to the radii values as opposed to the classical application of line fitting to position values. To quantify linearity, we assumed that each segment could be in one of three types: increasing in radii, decreasing in radii, or uniform in radii. The segments whose points had uniform radii corresponded to linear objects.

For each segment, an incremental line fitting algorithm was applied. Given a set of n points and their radii $r_t, t = 1, \dots, n$, a straight line was modeled with the function $r = at + b$ where a and b were the line parameters. The measure of how well a line fitted a set of n points was computed using the least-squares error (LSE) criterion

$$LSE = \sum_{t=1}^n (at + b - r_t)^2 \quad (3)$$

where $at + b - r_t$ was the algebraic distance between the t 'th point and the line. The line parameters a and b that minimized Equation (3) were found by taking partial derivatives and solving for the unknowns as

$$\begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} \sum_{t=1}^n t^2 & \sum_{t=1}^n t \\ \sum_{t=1}^n t & \sum_{t=1}^n 1 \end{bmatrix}^{-1} \begin{bmatrix} \sum_{t=1}^n tr_t \\ \sum_{t=1}^n r_t \end{bmatrix}. \quad (4)$$

The incremental line fitting algorithm initialized a subsegment with the first two points in the segment. The algorithm walked along the segment and fitted a line to runs of points along this segment. Given a subsegment, a line was fitted to the points on the subsegment and the next point on the segment. This point was added to the subsegment if the LSE was less than a threshold. Otherwise, a new subsegment was started like the initial subsegment. The subsegment was selected as corresponding to a linear structure (i.e., having uniform radii) if the corresponding slope (line parameter a) was close to 0 (i.e., having an absolute value less than a threshold). After the linear subsegments of the skeleton were obtained, the next problem was to extract the image objects that corresponded to these subsegments. This problem was solved by iteratively thickening each subsegment by constraining the result with the thickening of the rest of the skeleton and the original candidate object areas. The resulting area for each subsegment was recorded as an accepted candidate object. The segment selection process is illustrated in Figure 5.

The final set of shape features consisted of the aspect (length/width) ratio for each resulting object. The length was calculated as the number of points on the skeleton of the corresponding subsegment, and the width was calculated as the average diameter for the points on the skeleton of the subsegment. The final decision for accepting a segment as a target object was done using a threshold on aspect ratio. Example detection results are shown in Figures 4(c) and 4(f).

VI. PERFORMANCE EVALUATION

The overall algorithm for the detection of linear strips of woody vegetation is summarized in Algorithm 1. Performance evaluation of woody vs. non-woody classification for the identification of candidate objects was presented in Section IV. In this section, we describe the object-based quantitative performance measures used for evaluating the accuracy of target detection with respect to different parameter settings, and present quantitative and qualitative results.




Fig. 5. Illustration of the iterative least-squares line fitting based segment selection process. The first column shows woody classification maps with one of the skeleton segments overlaid. The second column shows plots of radii of the points on these segments. The x -axis shows the points (t) and the y -axis shows the corresponding radii ($r_t, t = 1, \dots, n$). The subsegments corresponding to the lines found are shown with different colors. The red subsegments are accepted as linear objects as their slopes are close to 0 (i.e., having uniform radii).

A. Object-based performance measures

Quantitative evaluation of thematic classification accuracy has been well-studied in the literature using measures such as error rates computed from confusion matrices or the Kappa coefficient. However, such measures alone are not always sufficient indicators of the geometric accuracy of object detection [22]. Methods for assessing the geometric accuracy of image segmentation algorithms [23] are also available in the literature. However, the measures that are based on matches between two complete partitionings of the whole image are not directly applicable to the problem studied in this paper where the goal is to detect particular objects, not to partition the whole land cover. Furthermore, due to the absence of a rigorous definition of a hedge object and the fact that the delimitation of the ground truth remains approximate because of the limitations of the CAPI-based hedge detection, some geometric error measures such as the border (edge) error or the shape error (e.g., eccentricity) [22] are not always suitable indicators of the performance for this problem.

The object-based performance criteria used in this paper were adapted from the work of Hoover *et al.* [24] on the evaluation of range image segmentation algorithms. Hoover *et al.* [24] classified every pair of reference and output objects as correct detections, over-detections, under-detections, missed

Algorithm 1 Hedge Detection

Perform pixel level feature extraction
 Threshold NDVI {*parameter*: threshold}
 Classify remaining pixels as woody or non-woody
 Eliminate small noise components and compute candidate object map
 Compute skeleton and corresponding radii for candidate objects
 Prune skeleton by removing branches shorter than a threshold {*parameter*: length threshold}
 Perform top-hat transform and eliminate structures wider or narrower than an acceptable hedge {*parameters*: max-width and min-width values corresponding to diameters of disk structuring elements}
 Find linked point lists (segments) identified by junctions and end points
for all skeleton segments **do**
 Perform line fitting and find subsegments that are linear enough {*parameters*: least-squares fitting error threshold and slope threshold}
end for
 Find objects corresponding to selected subsegments
for all objects **do**
 Compute aspect (length/width) ratio
 Threshold based on aspect ratio {*parameter*: aspect threshold}
end for

detections or false alarms with respect to a threshold on the amount of overlap, in terms of the number of pixels, between these objects. Due to the approximations in the ground truth in this study, we performed the evaluation using skeletons. Since the skeleton is a good indicator of the geometric properties and the shape of a linear object, the match between a reference object and an output object was measured in terms of the overlap between their skeletons where a small dilation of each skeleton was used as a buffer for computing the overlap. Note that the criteria below are defined in terms of the overlap between two skeletons but can as well be computed using overlaps between objects to evaluate any object detection algorithm.

The input to the evaluation procedure includes:

- Objects in the ground truth: O_i^{GT} , $i = 1, \dots, M$,
- Length of the skeletons of O_i^{GT} : L_i^{GT} , $i = 1, \dots, M$,
- Objects in the algorithm output: O_j^{AO} , $j = 1, \dots, N$,
- Length of the skeletons of O_j^{AO} : L_j^{AO} , $j = 1, \dots, N$.

First, we construct an $M \times N$ table where each entry C_{ij} corresponds to the length of the overlap between the skeletons of objects O_i^{GT} and O_j^{AO} . If there is no overlap between the skeletons of the two objects, $C_{ij} = 0$. If they share the same skeleton, $C_{ij} = L_i^{\text{GT}} = L_j^{\text{AO}}$. Then, given a threshold T , the results can be classified into five types of detections:

- 1) *Correct detection*: A pair of objects O_i^{GT} and O_j^{AO} is classified as an instance of correct detection if
 - $C_{ij} \geq T \times L_j^{\text{AO}}$ (at least T percent of the skeleton of O_j^{AO} overlaps with the skeleton of O_i^{GT} with an

- overlap score of $f_1 = C_{ij}/L_j^{\text{AO}}$), and
 - $C_{ij} \geq T \times L_i^{\text{GT}}$ (at least T percent of the skeleton of O_i^{GT} overlaps with the skeleton of O_j^{AO} with an overlap score of $f_2 = C_{ij}/L_i^{\text{GT}}$).
- 2) *Over-detection*: An object O_i^{GT} and a set of objects $O_{j_1}^{\text{AO}}, \dots, O_{j_n}^{\text{AO}}$, $2 \leq n \leq N$, are classified as an instance of over-detection if
 - $C_{ij_k} \geq T \times L_{j_k}^{\text{AO}}$, $\forall k \in \{1, \dots, n\}$ (at least T percent of the skeleton of each $O_{j_k}^{\text{AO}}$ overlaps with the skeleton of O_i^{GT} with a total overlap score of $f_1 = \sum_{k=1}^n C_{ij_k} / \sum_{k=1}^n L_{j_k}^{\text{AO}}$), and
 - $\sum_{k=1}^n C_{ij_k} \geq T \times L_i^{\text{GT}}$ (at least T percent of the skeleton of O_i^{GT} overlaps with the union of the skeletons of $O_{j_1}^{\text{AO}}, \dots, O_{j_n}^{\text{AO}}$ with a total overlap score of $f_2 = \sum_{k=1}^n C_{ij_k} / L_i^{\text{GT}}$).
- 3) *Under-detection*: A set of objects $O_{i_1}^{\text{GT}}, \dots, O_{i_m}^{\text{GT}}$, $2 \leq m \leq M$, and an object O_j^{AO} are classified as an instance of under-detection if
 - $\sum_{k=1}^m C_{i_k j} \geq T \times L_j^{\text{AO}}$ (at least T percent of the skeleton of O_j^{AO} overlaps with the union of the skeletons of $O_{i_1}^{\text{GT}}, \dots, O_{i_m}^{\text{GT}}$ with a total overlap score of $f_1 = \sum_{k=1}^m C_{i_k j} / L_j^{\text{AO}}$), and
 - $C_{i_k j} \geq T \times L_{i_k}^{\text{GT}}$, $\forall k \in \{1, \dots, m\}$ (at least T percent of the skeleton of each $O_{i_k}^{\text{GT}}$ overlaps with the skeleton of O_j^{AO} with a total overlap score of $f_2 = \sum_{k=1}^m C_{i_k j} / \sum_{k=1}^m L_{i_k}^{\text{GT}}$).
- 4) *Missed detection*: An object O_i^{GT} that is not in any instance of correct detection, over-detection or under-detection is classified as missed detection.
- 5) *False alarm*: An object O_j^{AO} that is not in any instance of correct detection, over-detection or under-detection is classified as false alarm.

Although these definitions result in a classification for every object in the ground truth and the algorithm output, these classifications may not be unique for $T < 1.0$ as discussed in [24]. However, for $0.5 < T < 1.0$, any object can contribute to at most three classifications (at most one correct detection, one over-detection and one under-detection). If an object is included only in a single classification instance (of correct detection, over-detection or under-detection), that instance is used as its unique classification. When an object participates in two or three classification instances, the instance with the highest score is selected for that object. The score for a classification instance is computed as the average of the two overlap scores $((f_1 + f_2)/2)$ in the corresponding definition. For equal scores, we bias toward selecting correct detection, then over-detection, then under-detection. An overlap threshold of $T = 0.6$ was used in the experiments in this paper.

Precision and recall have been commonly used in the literature to measure how well the detected objects correspond to the ground truth objects [6]. Recall can be interpreted as the number of true positive objects detected by the algorithm, while precision evaluates the tendency of the algorithm for false positives. Once all reference and output objects are classified into instances of correct detections, over-detections, under-detections, missed detections or false alarms, precision

TABLE III

THE PARAMETER SETTINGS USED IN THE TARGET DETECTION EXPERIMENTS. THE PARAMETERS ARE DESCRIBED IN ALGORITHM 1 AND IN THE TEXT.

	Values
Length (pixels)	25, 40, 50
Max-width (pixels)	50, 80, 100
LSE	0.3, 0.4, 0.5
Slope	0.1, 0.15, 0.2
Aspect	2, 3, 4

and recall are computed as

$$\text{precision} = \frac{\# \text{ of correctly detected objects}}{\# \text{ of all detected objects}} = \frac{N - FA}{N} \quad (5)$$

$$\text{recall} = \frac{\# \text{ of correctly detected objects}}{\# \text{ of all objects in the ground truth}} = \frac{M - MD}{M} \quad (6)$$

where FA and MD are the number of false alarms and missed detections, respectively. Finally, the F_β measure that provides a way of combining precision and recall into a single measure that falls between the two is computed as

$$F_\beta = \frac{(\beta^2 + 1) \times \text{precision} \times \text{recall}}{\beta^2 \times \text{precision} + \text{recall}}. \quad (7)$$

The F_β measure attaches β times as much importance to recall as precision. The F_2 measure ($\beta = 2$) was used in the experiments below to rank the performances of different parameter settings.

B. Results and discussion

The object-based performance criteria described above were used to evaluate the cross-landscape and site specific performances of the woody vegetation classifier and the shape-based target detection algorithm with respect to different parameter settings. Similar to Section IV, 4 training scenarios were considered:

- training the woody vegetation classifier and parameter selection for the shape-based target detection algorithm using samples from all sites, and
- training the woody vegetation classifier and parameter selection for the shape-based target detection algorithm using samples only from individual sites (3 scenarios).

The woody vegetation classifiers used were the ones described in Section IV and evaluated in Tables I and II. Note that the shape-based target detection step did not need any training, and was the same for all scenarios.

We considered 3 different values for each of the length threshold for pruning the skeleton, the max-width parameter for the top-hat transform, the least-squares fitting error (LSE) threshold, the slope threshold, and the aspect threshold. These values are shown in Table III. These settings corresponded to $4 \times 3^5 = 972$ parameter combinations for 4 scenarios. The NDVI threshold was fixed at 0.3 and the min-width parameter for the top-hat transform was fixed at 5 pixels (corresponding to 3 meters at 60 cm spatial resolution).

Table IV summarizes the parameter settings that obtained the best performance among all combinations. When all parameter combinations were considered, the following conclusions can be derived:

- The site specific training and parameter selection resulted in higher overall accuracies than the case where samples from all sites were used together. The former achieved an overall precision of 0.3523 and an overall recall of 0.5869 (the average of the last three rows in Table IV) compared to the 0.3212 precision and 0.5812 recall of the latter (the first row of Table IV). This was an expected result as site specific training of the woody vegetation classifier also performed better in the experiments in Section IV. However, the cross-landscape performance of the latter can also be considered acceptable, given the complexity of the targets of interest in different landscapes.
- When the performances for individual sites were considered, site specific training and parameter selection significantly improved the accuracy for *Paphos* sites (15.47% increase in the F_2 measure) but caused a slight decrease in the accuracies for *Baden* and *Decin* sites (1.26% and 3.75% decrease, respectively, in the F_2 measure). The improved accuracy for *Paphos* can be attributed to better training and parameter selection for this data set's hedges with much different characteristics than the hedges in other data sets. On the other hand, the hedges in the *Baden* and *Decin* sites were more similar to each other so they benefited from joint training when samples from all sites were used together. For all settings, the *Baden* and *Decin* sites received higher accuracies than the *Paphos* sites. This was also expected as the *Paphos* sites represent a rather extreme situation of thin hedges in a very fragmented environment containing many other small linear features.
- The length threshold for pruning the skeleton was selected as 25 for *Baden* and 50 for *Decin* and *Paphos* when samples from individual sites were used. Since the boundaries of *Baden* hedges were typically smoother than the boundaries of hedges in *Decin* and *Paphos* sites, a length threshold of 25 was sufficient for pruning for the former. The hedges in *Decin* sites were typically much wider than the hedges in other sites so a larger threshold of 50 was needed for pruning the skeletons and reducing the false alarms. The *Paphos* hedges were the most curly among all so a threshold of 50 was selected for pruning and smoothing. Furthermore, the *Paphos* sites contained a lot of small and thin structures in the woody vegetation map so the selected length threshold also helped eliminating some of these false alarms. The length threshold was selected as 40 as a tradeoff when samples from all sites were used.
- The best performing max-width threshold varied among different sites. The selection of 50, 80, and 100 as the best parameters for *Paphos*, *Baden*, and *Decin*, respectively, using samples from individual sites reflected the relative sizes and the expected width of acceptable hedges specified by the experts in these sites very well. These values corresponded to maximum expected width of 30 m, 48 m, and 60 m, respectively, at 60 cm spatial resolution. The max-width value of 80 was selected as a tradeoff when samples from all sites were used.
- The best LSE parameter also varied among different sites.

TABLE IV

THE PARAMETER SETTINGS THAT OBTAINED THE BEST PERFORMANCE AMONG ALL 972 COMBINATIONS. THE TRAINING COLUMN CORRESPONDS TO THE 4 TRAINING AND PARAMETER SELECTION SCENARIOS DESCRIBED IN THE TEXT. THE VALIDATION COLUMN CORRESPONDS TO THE SOURCE OF THE HEDGE OBJECT SAMPLES USED TO COMPUTE THE PERFORMANCE MEASURES IN EACH ROW. TD: TOTAL NUMBER OF OBJECTS IN THE ALGORITHM OUTPUT. GT: TOTAL NUMBER OF GROUND TRUTH HEDGE OBJECTS IN THE VALIDATION SITES. THE PARAMETERS AND THE PERFORMANCE MEASURES ARE DEFINED IN THE TEXT.

Training	Validation	Length	Max-width	LSE	Slope	Aspect	TD	Precision	GT	Recall	F ₂
All	All	40	80	0.3	0.2	2	688	0.3212	351	0.5812	0.5002
All	Baden	40	80	0.3	0.2	2	203	0.3990	131	0.6031	0.5471
All	Decin	40	80	0.3	0.2	2	191	0.3822	95	0.6421	0.5652
All	Paphos	40	80	0.3	0.2	2	294	0.2279	125	0.5120	0.4098
Baden	Baden	25	80	0.4	0.2	3	236	0.3814	131	0.6031	0.5402
Decin	Decin	50	100	0.3	0.2	2	154	0.3961	95	0.6000	0.5440
Paphos	Paphos	50	50	0.5	0.1	3	226	0.2920	125	0.5600	0.4732

A value of 0.3 was selected for *Decin* as the hedges in these sites were the least curly whereas a larger value of 0.5 was needed for *Paphos* hedges as these hedges had the most curly boundaries producing a larger variance in the corresponding width.

- The best performing slope parameter was selected as 0.2 for all scenarios except the site specific *Paphos* case. A smaller value of 0.1 was needed for the latter to eliminate the false alarms caused by small segments with highly curly boundaries.
- The aspect threshold was selected as 2 for the *Decin* hedges and 3 for the *Baden* and *Paphos* hedges because the hedges in the *Decin* sites were much wider than the hedges in the others. The aspect threshold of 3 was preferred for *Baden* and *Paphos* to eliminate false alarms from small segments. The parameter was selected as 2, which was the least restrictive setting, when samples from all sites were used.

Example detections are shown in Figure 6. The visual interpretation of the results showed that recall was actually higher than what was reported in Table IV. For example, many ground truth hedges had corresponding detected skeleton segments in Figures 6(e)–6(h). However, some of these were not counted as correct detections because of the limitations of the 60% overlap requirement ($T = 0.6$) between the skeletons. This requirement could not always be satisfied due to some shifts in the skeletons because the algorithm output was affected by the shadows within and on the sides of the woody vegetation but the ground truth did not include these shadows. Similarly, precision was also observed to be higher than the values in Table IV. Due to the absence of a rigorous definition of a hedge object and the fact that the delimitation of the ground truth remains approximate because of the limitations of the CAPI-based hedge detection, the images contained several linear structures that could be interpreted as woody vegetation but were not included in the ground truth. This was a particular problem for the *Paphos* sites as several objects that appeared in the algorithm output and could be considered as linear woody vegetation were counted as false alarms (Figures 6(g)–6(h)).

When the overall detections were considered, the following sources of error were identified. Some of the missed detections were caused by the errors during the identification of candidate objects. The pixel-based classification of woody vegetation sometimes could not isolate the hedges from the nearby

vegetation, and this caused the following shape-based target detection step to fail to extract correct shape and structure information for these hedge objects. Some of the false alarms were caused by groups of individual but nearby trees in orchards. This was observed to be a problem in *Paphos* data when samples from all sites were used for training, but significantly decreased with site specific training as smaller scale texture features gained more weight in the classifier for these images. Some other false alarms were caused by groups of trees in residential areas (Figures 6(e)–6(f)). We believe that these errors can be reduced by incorporating a contextual decision step, for example, by requiring a detected object to be neighboring an agricultural field. Yet another source of false alarm was the linear vegetation that did not look woody enough and was not included in the ground truth. Such linear vegetation structures sometimes gave high responses to the texture features, and appeared among the candidate objects that were tested for linearity (Figures 6(g)–6(h)). Finally, the difficulty in defining a hedge object and the associated vagueness in the connectedness of trees in a hedge structure led to missed detection of several small segments that could be considered as part of a hedge structure but were not fully connected to that structure. Using the same thresholds for all segments rejected some of these short segments. A possible solution to this problem may be the use of hysteresis thresholding like commonly done in edge detection for not eliminating short linear segments if they are near long linear segments. Given the difficulty in the definition of hedge objects and the large amount of variation in their appearance in different sites, we can conclude that the proposed framework that exploited the spectral, textural, and shape information using hierarchical feature extraction and decision making steps was successful in the detection of such objects in a wide range of landscapes.

VII. CONCLUSIONS

We described a framework for automatic detection of linear strips of woody vegetation in agricultural landscapes. The detection of these objects was considered important because they belong to the farm property that constituted an element of the landscape’s ecological infrastructure, and the monitoring of such structures is important for evaluating the environmental impact of the agriculture sector.

Detection of such specific target objects (hedges) necessitated a multi-scale and multi-feature strategy as no single




Fig. 6. Example results for hedge detection. The first row shows the object level ground truth. The second row shows the segment classification according to the object-based performance measures. The ground truth objects that are correctly detected (red), over-detected (green), under-detected (blue), and missed (gray) are shown as regions. The algorithm outputs that correspond to correct detections (red), over-detections (green), under-detections (blue), and false alarms (gray) are shown as overlaid skeleton segments. The third row shows the detected objects overlaid on the visible bands.

feature could achieve good localization performance individually. Multi-scale texture features could characterize the fine texture of individual trees as well as the coarse texture of their different groupings. Feature selection experiments showed that a Gaussian classifier that used a combination of spectral and textural features could identify the woody vegetation areas as candidate objects. An important step was a novel modeling of shape information as the objects of interest were often connected to other larger groups of trees, and often followed natural boundaries that did not necessarily exhibit a perfectly straight structure. The model involved morphological top-hat transforms to locate the woody vegetation areas that fell within the width limits of an acceptable hedge, and a skeletonization and iterative least-squares procedure that quantified the linearity of the objects. The proposed solution was generic enough for adaption to the detection of other linear object classes (e.g., rivers, roads, paths) with natural boundaries. Extensive

experiments using QuickBird imagery from three European Union member states showed that the proposed algorithms provided good localization of the target objects in landscapes with very different characteristics.

REFERENCES

- [1] S. Aksoy, H. G. Akcay, R. G. Cinbis, and T. Wassenaar, "Automatic mapping of linear woody vegetation features in agricultural landscapes," in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, Boston, Massachusetts, July 6–11, 2008.
- [2] DEFRA, "The hedgerows regulations 1997: A guide to the law and good practice," Department for Environment, Food and Rural Affairs, London, UK, Tech. Rep., 2002.
- [3] J. Baudry, R. G. H. Bunce, and F. Burel, "Hedgerows: An international perspective on their origin, function and management," *Journal of Environmental Management*, vol. 60, no. 1, pp. 7–22, September 2000.
- [4] M. W. Thornton, P. M. Atkinson, and D. A. Holland, "Sub-pixel mapping of rural land cover objects from fine spatial resolution satellite sensor imagery using super-resolution pixel-swapping," *International Journal of Remote Sensing*, vol. 27, no. 3, pp. 473–491, February 2006.

- [5] RPA and DEFRA, "Single payment scheme: Cross compliance — guidance for the management of habitats and landscape features," Rural Payments Agency and Department for Environment, Food and Rural Affairs, London, UK, Tech. Rep., 2004.
- [6] H. G. Akcay and S. Aksoy, "Automatic detection of geospatial objects using multiple hierarchical segmentations," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 46, no. 7, pp. 2097–2111, July 2008.
- [7] G. Rondeaux, M. Steven, and F. Baret, "Optimization of soil-adjusted vegetation indices," *Remote Sensing of Environment*, vol. 55, no. 2, pp. 95–107, February 1996.
- [8] K. Johansen and S. Phinn, "Mapping indicators of riparian vegetation health using IKONOS and Landsat-7 ETM+ image data in Australian tropical savannas," in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, vol. 3, Anchorage, Alaska, September 20–24, 2004, pp. 1559–1562.
- [9] L. J. Quackenbush, "A review of techniques for extracting linear features from imagery," *Photogrammetric Engineering & Remote Sensing*, vol. 70, no. 12, pp. 1383–1392, December 2004.
- [10] F. Tupin, H. Maitre, J.-F. Mangin, J.-M. Nicolas, and E. Pechersky, "Detection of linear features in SAR images: Application to road network extraction," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, no. 2, pp. 434–453, March 1998.
- [11] H. Mayer, "Automatic object extraction from aerial imagery — a survey focusing on buildings," *Computer Vision and Image Understanding*, vol. 74, no. 2, pp. 138–149, May 1999.
- [12] V. Onana, E. Trouve, G. Mauris, J.-P. Rudant, and E. Tonye, "Linear features extraction in rain forest context from interferometric SAR images by fusion of coherence and amplitude information," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 11, pp. 2540–2556, November 2003.
- [13] M. W. Thornton, P. M. Atkinson, and D. A. Holland, "A linearized pixel-swapping method for mapping rural linear land cover features from fine spatial resolution remotely sensed imagery," *Computers & Geosciences*, vol. 33, pp. 1261–1272, 2007.
- [14] C. Thenail and J. Baudry, "Variation of farm spatial land use pattern according to the structure of the hedgerow network (bocage) landscape: A case study in northeast Brittany," *Agriculture, Ecosystems and Environment*, vol. 101, no. 1, pp. 53–72, January 2004.
- [15] M. Lennon, M. C. Mouchot, G. Mercier, and L. Hubert-Moy, "Segmentation of hedges on CASI hyperspectral images by data fusion from texture, spectral and shape analysis," in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, vol. 2, Honolulu, Hawaii, July 24–28, 2000, pp. 825–827.
- [16] C. Vannier and L. Hubert-Moy, "Detection of wooded hedgerows in high resolution satellite images using an object-oriented method," in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, Boston, Massachusetts, July 6–11, 2008.
- [17] K. Tansey, I. Chambers, A. Anstee, A. Denniss, and A. Lamb, "Object-oriented classification of very high resolution airborne imagery for the extraction of hedgerows and field margin cover in agricultural areas," *Applied Geography*, vol. 29, no. 2, pp. 145–157, April 2009.
- [18] A. K. Shackelford and C. H. Davis, "A hierarchical fuzzy classification approach for high-resolution multispectral data over urban areas," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 9, pp. 1920–1932, September 2003.
- [19] B. S. Manjunath and W. Y. Ma, "Texture features for browsing and retrieval of image data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 837–842, August 1996.
- [20] P. Soille, *Morphological Image Analysis*, 2nd ed. Springer, 2002.
- [21] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. John Wiley & Sons, Inc., 2000.
- [22] L. Bruzzone and C. Persello, "A novel protocol for accuracy assessment in classification of very high resolution multispectral and SAR images," in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, Boston, Massachusetts, July 6–11, 2008.
- [23] A. Ortiz and G. Oliver, "On the use of the overlapping area matrix for image segmentation evaluation: A survey and new performance measures," *Pattern Recognition Letters*, vol. 27, no. 16, pp. 1916–1926, December 2006.
- [24] A. Hoover, G. Jean-Baptiste, X. Jiang, P. J. Flynn, H. Bunke, D. B. Goldgof, K. Bowyer, D. W. Eggert, A. Fitzgibbon, and R. B. Fisher, "An experimental comparison of range image segmentation algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 7, pp. 673–689, July 1996.



Selim Aksoy (S'96-M'01) received the B.S. degree from the Middle East Technical University, Ankara, Turkey, in 1996 and the M.S. and Ph.D. degrees from the University of Washington, Seattle, in 1998 and 2001, respectively.

He has been an Assistant Professor at the Department of Computer Engineering, Bilkent University, Ankara, since 2004, where he is also the Co-Director of the RETINA Vision and Learning Group. Before joining Bilkent, he was a Research Scientist at Insightful Corporation, Seattle, where he was involved in image understanding and data mining research sponsored by the National Aeronautics and Space Administration, the U.S. Army, and the National Institutes of Health. During 1996–2001, he was a Research Assistant at the University of Washington, where he developed algorithms for content-based image retrieval, statistical pattern recognition, object recognition, graph-theoretic clustering, user relevance feedback, and mathematical morphology. During the summers of 1998 and 1999, he was a Visiting Researcher at the Tampere International Center for Signal Processing, Tampere, Finland, collaborating in a content-based multimedia retrieval project. His research interests include computer vision, statistical and structural pattern recognition, machine learning and data mining with applications to remote sensing, medical imaging, and multimedia data analysis.

Dr. Aksoy is a member of the IEEE Geoscience and Remote Sensing Society, the IEEE Computer Society, and the International Association for Pattern Recognition (IAPR). His received the CAREER Award from the Scientific and Technological Research Council of Turkey (TUBITAK) in 2004. He was one of the Guest Editors of the special issues on Pattern Recognition in Remote Sensing of IEEE Transactions on Geoscience and Remote Sensing and Pattern Recognition Letters in 2007 and 2009, respectively. He served as the Vice Chair of the IAPR Technical Committee 7 on Remote Sensing during 2004–2006, and is currently the Chair of the same committee for 2006–2010. He is also an Associate Editor of Pattern Recognition Letters.



H. Gökhan Akçay received the B.S. and M.S. degrees in computer engineering from Bilkent University, Ankara, Turkey, in 2004 and 2007, respectively. He is currently pursuing the Ph.D. degree in computer engineering at Bilkent University.

Since September 2004, he has been a Teaching and Research Assistant at the Department of Computer Engineering, Bilkent University. His research interests include statistical and structural pattern recognition, computer vision and remote sensing image analysis.



Tom Wassenaar received the M.Sc. degree in physical geography from the University of Amsterdam, The Netherlands, in 1995. As an EU funded Marie Curie research fellow with the French agricultural research institute INRA he received his Ph.D. degree in soil science from the French École Nationale Supérieure d'Agronomie in Montpellier, in 2001. His thesis focused on the automatic processing of very high resolution (aerial) imagery for the detection and mapping of parameters relevant for distributed hydrological modelling in a Mediterranean vineyard environment. He has since been performing research on agri-environmental issues, particularly with the Food and Agricultural Organization of the UN. Currently with the Monitoring Agricultural Resources unit of the Institute for the Protection and Security of the Citizen at the Joint Research Center of the European Commission, he continues in parallel to develop agriculture related very high resolution remote sensing applications.